

Rational Verification in Iterated Electric Boolean Games*

Youssef Oualhadj

LACL, U-PEC
Paris, France

youssef.oualhadj@lACL.fr

Nicolas Troquard

LACL, U-PEC
Paris, France

nicolas.troquard@lACL.fr

Electric boolean games are compact representations of games where the players have qualitative objectives described by LTL formulae and have limited resources. We study the complexity of several decision problems related to the analysis of rationality in electric boolean games with LTL objectives. In particular, we report that the problem of deciding whether a profile is a Nash equilibrium in an iterated electric boolean game is no harder than in iterated boolean games without resource bounds. We show that it is a PSPACE-complete problem. As a corollary, we obtain that both rational elimination and rational construction of Nash equilibria by a supervising authority are PSPACE-complete problems.

1 Introduction

We study multiagent systems populated with self-interested agents who interact repeatedly and are limited in their actions by a limited amount of energy. We investigate the computational aspects of deciding whether a collective, non-cooperative, behaviour is rational.

Electric boolean games The formalism under consideration was introduced in the second part of [19] but the decision problems were left open. They extend naturally the models of multi-player boolean games [7], one-shot electric games [19], and iterated boolean games [17]. Boolean games have occupied an important position in the recent formal AI literature. This line of work is an effort in formalisation of game theoretical situations with boolean games (see previously cited work and e.g., [24, 15]).

Strategically, the players in Iterated Electric Boolean Games (Sec. 2) are intricately mixing qualitative and quantitative considerations. Not only do they need to find a strategy that helps them satisfy their qualitative objective over time, they need to do so, seeking to keep the interaction alive so as not to run out of energy and fail to be able to perform a single action. This can be illustrated by the next simple example.

Example 1. *Isabella and Jules are two demanding kids. Isabella's objective towards happiness is to be granted a new comic book on a regular basis, and Jules' objective is to be granted a new jigsaw puzzle just as often. Their mom's objective is naturally to have all requests eventually fulfilled. Whether they ask for a new item or not, it costs zero to the kids either way. They never incur any costs. Buying a new comic book however, will cost \$4 to their mother, and getting a new jigsaw puzzle will cost her \$6. Each day, each item that is not bought will earn Mom \$1. Isabella and Jules, being what they are, decide that their behaviour to satisfy their objective is to ask a new item all the time. Fortunately, Mom is going to cope with it by waiting 5 days, buying a new comic book and a new jigsaw puzzle on the 6-th day, and repeating. It results in a collective behaviour which is rational as we shall explain later on.*

*The first author is partially supported by the ANR project EQINOCS (ANR-11-BS02-004)

Boolean games as compact game representations Solving problems on an input only makes sense when the input is reasonable. Possible worlds and relational semantics are commonly used to model multi-agent systems. However, describing a complex system in terms of possible worlds is often impractical. In fact, the size of the description of a system as a transition system typically grows exponentially in the number of variables in the system. For instance, model checkers for Alternating-time Temporal Logic make use of Reactive Modules [2] or Interpreted Systems [21] to overcome the difficulty. The powers of agents and coalitions are derived from the ability to control the value of some variables, thus bringing about some change to the system. Boolean games [18, 8] are such compact representations which in addition also integrate agents' preferences. They recently have been widely used to study various phenomena relevant to artificial intelligence [15, 6, 5, 16, 24].

Boolean Games are multi-player games where each player controls a set of propositional variables and has a qualitative preference represented by a propositional formula over the set of variables in the system. An action for a player is to assign a valuation to the propositional variables she controls. Iterated Boolean Games [17] are a variant of Boolean games where the players repeat the interaction infinitely often, and where their qualitative objectives are represented as LTL formulas over the set of variables in the system.

Electric Boolean Games [19] are an extension of Boolean Games where agents are assigned an initial energy endowment and taking actions has a cost, positive or negative. Already in [19], the authors define an iterated version of Electric Boolean Games, but they do not investigate their strategic aspects.

Design of safe computer systems In theoretical computer science, and particularly in the design and verification of computer systems, two-player zero-sum games have been extensively studied and used with great success [3, 23]. Recently, researchers have brought their attention to introducing quantitative restrictions for the players. For instance games where the system has to accomplish a task while maintaining its resource level above zero was modelled using Mean payoff Parity games [13], or Energy Parity Games [12]. This line of work was naturally extended by the study of the so-called multi-objective games with actual implementation [9]. In a multi objective game, a protagonist player wants to achieve a conjunction of goals, and the antagonist player wants to achieve the exact opposite. Nevertheless, the pessimistic assumption that a system and its environment always have opposite interests is not always realistic. Therefore, multiplayer games seem to be a more suitable formalism [10]. Indeed, the environment is considered to be another player with her own goal. In order to study those games, the solution concept of choice was Nash equilibria as it is a sensible formalisation of rationality [11]. In an electric boolean game, each agent has to partake in a cooperation that keeps the system alive. Namely, every single player has to make sure that none of the other players is running out of resource. This approach can be seen as an intermediate setting between non-cooperative and cooperative games. Actually, this can also be seen as a new definition of multi-objective games in the setting of multi-player games; Every player has a personal goal with no incentive to cooperate and second goal where it is best for her to cooperate.

Engineering multiagent systems Some plays of a game may appear better than others by some supervising authority. Some strategic equilibria in a game may be undesirable, while play which are not equilibria might be seen as desirable. A supervising authority could have the power to redistribute the resources available in the system so as to achieve better equilibria from their point of view. Dealing with resources such as energy, it then becomes interesting to study how much different the game would be, were the endowments of the players be different. As in [19], it is very natural to consider resource

redistributions that allow one to eliminate ‘bad’ equilibria and/or construct ‘good’ equilibria.

Apart from [24] and [19], looking into ways of engineering a game’s outcome has also been considered in [1]. The authors propose a framework where the winning conditions can be modified at a cost, thus changing the strategic equilibria of the game.

Contributions Our main result is the PSPACE membership for rational verification i.e., given a strategy profile decide whether it is a Nash equilibrium (Sec. 3). Note that the computational complexity in the electric case matches the one in the non-electric case. Our proof differs from the one in [17] for the non-electric case. Indeed, a straightforward adaptation of their proof would fail for it relies on a translation of the input into a well chosen LTL formula. In the electric case, one has to pay particular attention to the electric constraints (c.f., Ex. 7). This is a quantitative ingredient that is absent from LTL. We overcome this difficulty as follows. We construct a one-player game played on a weighted graph. This allows us to encode the behaviour of the possible deviator together with the electric constraints in an existing formalism, viz., *Energy Büchi games* [14]. We prove that a rational deviation exists iff this one-player game contains a winning strategy. The size of the constructed one-player game may be exponential in the size of the input. However, on-the-fly automata-theoretic techniques allow one to maintain a PSPACE upper-bound for the problem of finding a winning strategy. Finally, to decide in PSPACE whether a strategy profile is a Nash equilibrium, it suffices to guess a deviator and check whether she has a winning strategy in her one-player game.

Solving rational verification facilitates the access to more problems. We show (Sec. 4) that the problems of resource redistribution come out as corollaries. We leave open the more challenging problem of rational synthesis for which rational verification is a stepping stone; Rational verification is to model checking what rational synthesis is to model synthesis.

A full version is available in [22].

2 Iterated Electric Boolean Games

Definition 2 (Electric Boolean Games). *An electric boolean game (EBG for short) is a tuple $\mathcal{B} = (N, A, \Phi, c, e)$ where: $N = \{1, \dots, n\}$ is a finite set of players. $A = \cup_{i=1}^n A_i$ with A_i are the atoms controlled by player i and (A_1, \dots, A_n) forms a partition of A . $\Phi = \{\phi_1, \dots, \phi_n\}$ where ϕ_i is the objective of player i . $c : A \times \{\perp, \top\} \rightarrow \mathbb{Z}$ is a cost function. $e : N \rightarrow \mathbb{N}$ is an endowment function.*

We denote \mathcal{T} the set $\{\perp, \top\}$ and for any set E , \mathcal{T}^E the set of mappings from E to \mathcal{T} , the set of all the finite sequences over E is E^* , and E^ω is the set of all the infinite sequences over E .

Let X be a set of atomic propositions, a valuation of X is a total function $v \in \mathcal{T}^X$. The cost of a valuation v is given by $\text{cst}(v) = \sum_{p \in X} c(p, v(p))$. An *action* of player i is to assign a valuation to each variable in the set A_i of the atoms she controls.

We consider the setting of concurrent and infinitely repeated electric boolean games, where players choose their actions simultaneously and for an infinite duration. We consider objectives in Φ which are specified by LTL formulas over the atoms of A ([4, Chap. 5]). Formulas of LTL are defined by the following grammar: $\phi ::= p \mid \phi \wedge \phi \mid \neg \phi \mid X\phi \mid \phi \cup \phi$ where $p \in A$. The other propositional operands and temporal operators (F, G) can be defined as usual.

We need to introduce some useful terminology to talk about repeated games and define the semantics of LTL formulas over $(\mathcal{T}^A)^\omega$.

A *history* in a repeated electric boolean game is a word in $(\mathcal{T}^A)^*$. That is, a finite sequence of valuations for the set A of boolean variables. A *play* is an infinite sequence in $(\mathcal{T}^A)^\omega$. Given a play ρ ,

we note $\rho[t]$ the t -th valuation function in ρ . We note $\rho[t\dots]$ the suffix of ρ starting at $\rho[t]$, and $\rho[\dots t]$ the prefix of ρ ending at $\rho[t]$ which is a history of size $t + 1$.

LTL objectives are evaluated over a play ρ of the game. For $p \in A$, and for ϕ and ψ two LTL formulas:

$$\begin{aligned} \rho \models p &\text{ iff } \rho[0](p) = \top & \rho \models \neg\phi &\text{ iff } \rho \not\models \phi \\ \rho \models X\phi &\text{ iff } \rho[1\dots] \models \phi & \rho \models \phi \wedge \psi &\text{ iff } \rho \models \phi \text{ and } \rho \models \psi \\ \rho \models \phi U \psi &\text{ iff } \exists i \geq 0, \rho[i\dots] \models \psi \text{ and } \forall 0 \leq j < i, \rho[j\dots] \models \phi \end{aligned}$$

The formula $X\phi$ holds true on ρ if ϕ is true next. The formula $\phi U \psi$ holds true on ρ if ϕ is true at least until ψ is true.

In order to play, the players choose their actions according to a strategy. A *strategy* for player i is a mapping that takes as input a history and outputs a valuation for each atom controlled by player i . Formally a strategy σ_i for player i is a mapping $\sigma_i : (\mathcal{T}^A)^* \rightarrow \mathcal{T}^{A_i}$. We note Σ_i the set of strategies of player i .

A *strategy profile* σ is a vector $(\sigma_1, \dots, \sigma_n)$ specifying one strategy σ_i for each player $i \in N$. Given a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and a strategy τ_i for player i , we note (τ_i, σ_{-i}) the strategy profile $(\sigma_1, \dots, \tau_i, \dots, \sigma_n)$. Each strategy profile induces a play, and since we consider pure strategies, there is one and only one such play consistent with σ . We denote $\langle \sigma \rangle$ the play induced by the profile σ . It is defined inductively as follows: if $p \in A_i$ then $\langle \sigma \rangle[0](p) = \sigma_i(\varepsilon)(p)$, and for $t \geq 0$, $\langle \sigma \rangle[t+1](p) = \sigma_i(\langle \sigma \rangle[\dots t])(p)$.

The endowment $e(i)$ of each player i specified in the definition of an electric boolean game, represents the initial resources of the player. While playing the game following a strategy, this endowment grows as the player takes an action of negative cost and shrinks as the player takes an action of positive cost.

We will say that the strategy profile σ is feasible in an iterated EBG if it does not over-consume the endowed resources, in the sense that, every player's strategy σ_i can be infinitely executed without ever causing the player's compound endowment to go under 0. We make it more formal.

Consider an EBG (N, A, Φ, c, e) and a strategy profile σ . The *compound endowment* of player i at the t -th step of the play $\langle \sigma \rangle$ is defined with $E_i^\sigma(0) = e(i)$, and

$$E_i^\sigma(t+1) = E_i^\sigma(t) - \text{cst}(\sigma_i(\langle \sigma \rangle[\dots t]))$$

Thus, the strategy profile σ is *feasible* iff for each player $i \in N$, and for all $t \geq 0$ we have $E_i^\sigma(t) \geq 0$. In the strategy profile σ , we say that τ_i is a *feasible deviation* for player i iff (τ_i, σ_{-i}) is a feasible strategy profile.

Once an objective ϕ_i and a strategy profile σ are fixed, the payoff of σ for player i is defined as follows:

$$\text{Payoff}_i(\sigma) = \begin{cases} 1 & \text{if } \sigma \text{ is feasible, and } \langle \sigma \rangle \models \phi_i, \\ 0 & \text{otherwise.} \end{cases}$$

In the strategy profile σ , we say that τ_i is a *rational deviation* for player i iff $\text{Payoff}_i((\tau_i, \sigma_{-i})) > \text{Payoff}_i(\sigma)$.

Example 3. We formalise the game of Example 1 and model a strategy for the three participants. Let $\mathcal{B}^{c,e}$ be an EBG (N, A, Φ, c, e) where $N = \{I, J, M\}$, $A_I = \{r_I\}$, $A_J = \{r_J\}$, $A_M = \{g_I, g_J\}$. Evaluated to \top , the atoms r_I , r_J , g_I , g_J , respectively represent the facts that Isabella asks for a comic book, Jules asks for a jigsaw puzzle, Mom buys a comic book, and Mom buys a jigsaw puzzle. The costs are given by $c(r_I, \top) =$

$c(r_I, \perp) = c(r_J, \top) = c(r_J, \perp) = 0$, and $c(g_I, \perp) = c(g_J, \perp) = -1$, $c(g_I, \top) = 4$, and $c(g_J, \top) = 6$. We suppose that $e(I) = e(J) = e(M) = 0$. The objectives are given as $\Phi_M = G((r_I \rightarrow F(g_I)) \wedge (r_J \rightarrow F(g_J)))$, $\Phi_I = GF(g_I)$, and $\Phi_J = GF(g_J)$. The strategies of the kids continuously asking a new item and of the Mom buying one comic book and one jigsaw puzzle every 6 days result in a strategy profile whose payoff is 1 for everyone.

The strategies suggested at the end of Example 3 are depicted in Figure 1. They are instances of what we call *finite memory strategies*. We formalise the class of finite memory strategies next.

Definition 4 (Finite memory strategy). *Let $i \in N$ be a player, a finite memory strategy σ_i for player i consists of a finite set M called the memory, an initial memory state m^{in} in M , a mapping $\sigma_i^u : M \times \mathcal{T}^A \rightarrow M$ called the update function, and a mapping $\sigma_i^c : M \rightarrow \mathcal{T}^A$ called the choice function.*

We say that $(\sigma_1, \dots, \sigma_n)$ is a *finite memory profile* if for every $i \in N$, σ_i is a finite memory strategy. For instance, in the strategy of Figure 1c, the set M is $\{0, 1, 2, 3, 4, 5\}$, the initial memory state is 0, the update function is the edge relation and the choice function is illustrated by labels next to vertices¹.

3 Nash Equilibria in Electric Boolean Games

In [19], the authors introduced iterated electric boolean games but did not study their strategic aspects. Hence no solution concept was defined. However, the concept of Nash equilibria is one of most natural concept in multiplayer games.

Definition 5 (Nash equilibrium). *Let $\mathcal{B}^{c,e}$ be an EBG and σ be a strategy profile. We say that σ is a Nash equilibrium iff the following holds:*

1. $\forall t \geq 0, \forall i \in N, E_i^\sigma(t) \geq 0$,
2. $\forall i \in N, \forall \tau_i \in \Sigma_i, \text{Payoff}_i((\tau_i, \sigma_{-i})) \leq \text{Payoff}_i(\sigma)$.

Using our terminology, σ is a Nash equilibrium in $\mathcal{B}^{c,e}$ if and only if it is feasible and there is no rational deviation for any player. We note $\text{NE}(\mathcal{B}^{c,e})$ the set of Nash equilibria in the game $\mathcal{B}^{c,e}$. For instance, the strategy profile depicted in Figure 1 is a Nash equilibrium in the game of Examples 1 and 3


Definition 6 (Nash Equilibrium Membership). *Let $\mathcal{B}^{c,e}$ be an electric boolean game, and σ be a finite memory strategy profile. The Nash Equilibrium Membership (NEM) problem asks whether $\sigma \in \text{NE}(\mathcal{B}^{c,e})$.*

In order to build intuition regarding deviations, consider the following example


Example 7. *Let $\mathcal{B}^{c,e}$ be the following two-player game,*

$$\begin{aligned} A_1 &= \{p\}, A_2 = \{q\} , \\ \phi_1 &\equiv G((q \rightarrow Xp) \wedge (\neg q \rightarrow X\neg p)), \phi_2 \equiv Gq , \\ c(p, \top) &= 1, c(p, \perp) = -1, c(q, \top) = c(q, \perp) = 0, e(1) = e(2) = 0 . \end{aligned}$$


¹We omit the labels on the edges to highlight that for each player the update function depends only on the current memory state.



(a) Isabella's strategy.



(b) Jules' strategy.



(c) Mom's strategy.

Figure 1: A finite memory profile seen as finite graphs.

Consider the following strategy σ_1 for player 1 that assigns \top to p iff \top was assigned to q the previous round. We also consider the strategy σ_2 for player 2 that always assigns \perp to q .

We argue that the profile (σ_1, σ_2) is a Nash equilibrium. Clearly (σ_1, σ_2) is feasible. Let us show that player 2 does not have a rational deviation. In order to increase her payoff, player 2 has to always assign \top to q , call this new strategy τ . However, the deviation τ is not feasible. Indeed, player 1 is still following σ_1 , we obtain

$$\begin{aligned} \sigma_1(\varepsilon)(p) &= \perp \text{ with } E_1^{(\sigma_1, \tau)}(1) = 1, \\ \sigma_1(\{(p, \perp), (q, \top)\})(p) &= \top \text{ with } E_1^{(\sigma_1, \tau)}(2) = 0, \\ \sigma_1(\{(p, \perp), (q, \top)\}\{(p, \top), (q, \top)\})(p) &= \top \text{ with } E_1^{(\sigma_1, \tau)}(3) = -1, \end{aligned}$$

showing that the compound endowment drops below 0 after the third round. The plays induced by the two profiles are depicted in Figure 2.

This example shows that in order to perform a rational deviation, a player has to check the endowment of all the players and not only her own. We are now ready to state the main theorem of this paper.

Theorem 8. *NEM is a PSPACE-complete problem. It is PSPACE-hard even when there is only one player.*

To prove the theorem, we exhibit two constructions, c.f. Construction 1, and Construction 2. The former allows one to check the feasibility of a profile, while the latter allows one to check the existence of a rational deviation.

In Section 3.1, and Section 3.2 we let $\mathcal{B}^{c,e}$ be an EBG, and σ be a finite memory profile. Let also $(M_i, m_i^{\text{in}}, \sigma_i^{\text{u}}, \sigma_i^{\text{e}})$ be the finite memory strategy of player i in the profile σ .




Figure 2: plays induced by the profiles (σ_1, σ_2) and (σ_1, τ) .

3.1 Checking feasibility in PSPACE

We say that G is a d -weighted graph if G is associated with a weight function $w : E \rightarrow \mathbb{Z}^d$. For a vertex u and a vector w_0 in \mathbb{N}^d , a subset C of V is a nonnegative reachable cycle from u if the following holds. (i) There exists v in $C = \{u_j \mid 1 \leq j \leq k\}$, and a path $u_0, \dots, u_l, \dots, u_k$ such that $u_0 = u$, $u_l = v$, and $u_k = v$. (ii) For all $0 \leq t \leq k-1$ we have $w_0 - \sum_{j=0}^t w(u_j, u_{j+1}) \geq \{0\}^d$, and $\sum_{j=l}^{k-1} w(u_j, u_{j+1}) \leq \{0\}^d$. Positive cycles are defined as expected.

In order to prove Proposition 10 we use the results of [20]. In particular, given a d -weighted graph G , we can detect a nonnegative reachable cycle in polynomial time in the size of G .²

Our approach consists in constructing a n -weighted graph $G[\sigma]$ from the finite memory profile σ . This is achieved by Construction 1. We show that $G[\sigma]$ contains such a cycle iff σ is feasible.

We start first by giving the details of how $G[\sigma]$ is obtained.

Construction 1. $G[\sigma]$ consists of a finite set of vertices V , an edge relation $E \subseteq V \times V$, and weight function $w : E \rightarrow \mathbb{Z}^n$. $G[\sigma]$ is obtained as follows:

²The result of [20] is to find 0-cycles. To find nonnegative cycles, it suffices to transform a weighted graph G into G' by adding a reflexive edge of weight -1 to every vertex. This is a polynomial transformation. G has a nonnegative cycle iff G' has a zero-cycle.

- The vertices are $V = \prod_{i \in N} M_i$.
- For $v \in V$ we denote v_i the i -th component of v . Let $(u, v) \in V \times V$ be a couple of vertices, (u, v) is an edge in E if for each $i \in N$ we have $\sigma_i^u(u_i, X) = v_i$ where $X = \bigcup_{j \in N} \sigma_j^e(u_j)$ is the complete valuation over A prescribed by the profile σ .
- Finally, for $(u, v) \in E$,

$$w(u, v) = \left(\text{cst}(\sigma_1^e(u_1)), \dots, \text{cst}(\sigma_n^e(u_n)) \right).$$

The following lemma states the key property of Construction 1.

Lemma 9. *The finite memory strategy profile σ is feasible iff $G[\sigma]$ has a nonnegative reachable cycle from $u^0 = (m_1^{\text{in}}, \dots, m_n^{\text{in}})$ with initial credit e .*

A consequence of the above lemma is

Proposition 10. *We can check in PSPACE whether σ is feasible.*

3.2 Checking the existence of rational deviation in PSPACE

Now that we can check whether a profile is feasible, we need to show how to check the existence of rational deviation for a player.

We recall that $\mathcal{B}^{c,e}$, σ , and $\sigma_i = (M_i, m_i^{\text{in}}, \sigma_i^u, \sigma_i^e)$ are still fixed.

We need to introduce some technical material. A Büchi automaton \mathcal{A} is a tuple $\mathcal{A} = (Q, q_0, A, \Delta, F)$ where the Q is a finite set of states, q_0 is an initial state, A is a finite alphabet, Δ is relation in $Q \times A \times Q$, and F is a subset of states called accepting. We say that an infinite word w is recognised by \mathcal{A} if there exists an infinite path ρ in \mathcal{A} labelled by w such that ρ visits states in F infinitely many times. We also say that ρ is a run induced by w on \mathcal{A} . We define $\mathcal{L}_{\mathcal{A}}$ as the set of words recognised by \mathcal{A} . The reason we need Büchi automata is their strong link with LTL. Indeed, any LTL formula ϕ , can be associated to a Büchi automaton accepting all its models. The following theorem formalises this idea.

Theorem 11. *Let ϕ be a LTL formula, there exists a Büchi automaton \mathcal{A}_{ϕ} accepting the language \mathcal{L}_{ϕ} consisting of all the models of ϕ .*

The other formalism is *one-player games*. Let $G = (V, E, W)$ be a graph with a set of vertices V , a set of edges $E \subseteq V \times V$, and winning objective $W \subseteq V^{\omega}$. Strategies for these games are formalised by the following mapping $V^*V \rightarrow V$. Let σ be a strategy for the player, and u_0 a vertex in V . The play ρ starting in u_0 and consistent with σ is obtained as follows: $\rho[0] = u_0$, and for all $i > 0$, $\sigma(\rho[\dots i])$. The player wins if the play ρ is in W . A strategy σ is winning for the player from u_0 if the play consistent with σ is in W . Finite memory strategies can be defined in a similar fashion as for EBGs. In this paper, we use the so-called multi-objective games. Those are games where the player has to fulfil a combination of objectives at once.

Büchi objectives. We choose a set $F \subseteq V$ of accepting vertices. The winning objective W is $(V^*F)^{\omega}$. We denote this winning objective Büchi.

Energy objectives. Let $d > 0$ be a natural, $w_0 \in \mathbb{N}^d$ be an initial vector, and $w : E \rightarrow \mathbb{Z}^d$ be an energy function. The winning objective is the set $\{u_0 u_1 \dots \in V^{\omega} \mid \forall k \geq i, w_0 - \sum_{i=0}^k w(u_i, u_{i+1}) \geq \{0\}^d\}$. We denote this winning objective Energy.

The winning objective we are interested in is EnergyBüchi defined by $\text{Büchi} \cap \text{Energy}$.

Roughly speaking, given a profile σ and a player i , we construct a one-player EnergyBüchi game $G[\sigma_{-i}]$. The purpose of this game is to contain a winning strategy iff a rational deviation exists. Moreover, the winning strategy in $G[\sigma_{-i}]$ will be the deviation that player i uses to increase her payoff. Let us explain how to construct the one-player game $G[\sigma_{-i}]$.

Construction 2. We note V the set of vertices in $G[\sigma_{-i}]$, E the edge relation defined over $V \times \mathcal{T}^A \times V$, and the weight function w is a mapping from $V \times \mathcal{T}^A \rightarrow \mathbb{Z}^n$.

Let $\mathcal{A}_i = (Q, \mathcal{T}^{A_{\phi_i}}, q_0, \Delta, F)$ be an automaton accepting the language \mathcal{L}_{ϕ_i} .

The graph $G[\sigma_{-i}]$ is obtained as follows:

- The vertices are $V = Q \times \prod_{j \in N \setminus \{i\}} M_j$.
- Let v be a vertex in V , for $j \in N \setminus \{i\}$, v_j refers to the j -th component of v and v_i is the projection over Q . For $(u, v) \in V \times V$, and for every valuation $X \in \mathcal{T}^A$ we have (u, X, v) in E if
 - i) there exists $Y \in \mathcal{T}^{A_{\phi_i}}$ such that $(u_i, Y, v_i) \in \Delta$ and $Y \subseteq X$,
 - ii) the set $Z = Y \cup \bigcup_{j \in N \setminus \{i\}} \sigma_j^{\mathcal{E}}(u_j) \subseteq X$ and is consistent over A_{ϕ_i} i.e.

$$\forall p \in A_{\phi_i}, (p, \top) \in Z \implies (p, \perp) \notin Z,$$
 - iii) for each $j \in N \setminus \{i\}$ we have $\sigma_j^{\mathcal{U}}(u_j, X) = v_j$.
- The weight function is given by $\text{cst}(\sigma_j^{\mathcal{E}}(u_j))$ for every dimension $j \in N \setminus \{i\}$ and by $\sum_{p \in A_i} c(p, X(p))$ for dimension i .
- Finally, a vertex $v \in V$ is accepting if $v_i \in F$.

The intuition behind this construction is as follows. If player i can deviate rationally, then necessarily the new profile satisfies ϕ_i . This is why we use automaton \mathcal{A}_{ϕ_i} whose language is exactly those words that satisfy ϕ_i . Also, since we consider only unilateral deviations, the actions leading to the satisfaction of ϕ_i have to be compatible with the choices of other players, that is σ_{-i} . This is ensured by ii). Item iii) is a synchronisation between the action of the other player and the deviation of player i .

Thanks to the following lemma, we show that Construction 2 meets the desired intuition.

Lemma 12. Let σ be a finite memory profile, and i be a player such that $\text{Payoff}_i(\sigma) = 0$ then, i has a rational deviation iff there exists a winning strategy in $G[\sigma_{-i}]$.

As a consequence we obtain the core property for the existence of our PSPACE algorithm.

Proposition 13. Let σ be a finite memory profile, and i be a player such that $\text{Payoff}_i(\sigma) = 0$. We can check whether i has a rational deviation in PSPACE.

3.3 Proof of Theorem 8

We recall Theorem 8

Theorem 8. NEM is a PSPACE-complete problem. It is PSPACE-hard even when there is only one player.

Proof. If the profile is not feasible, return “no”. Otherwise, guess a possible deviator i (among the players with null payoff) and check whether she has a winning strategy in $G[\sigma_{-i}]$. Return “no” iff she has a winning strategy. Lemma 9 and Lemma 12 justify the correctness. Proposition 10 and Proposition 13 justify the upper-bound complexity.

To establish the hardness, one needs to notice that any BG is an EBG with endowment $\{0\}^N$ and $c : A \times \mathcal{T} \rightarrow \{0\}$. Thus the PSPACE lower bound established in [17, Prop. 2] holds for EBGs with LTL specifications. Since the proof is a reduction from LTL satisfiability to one-player iterated boolean games, NEM is hard even when there is only one player. \square

4 Resource redistributions

Having characterised the complexity of the problem of deciding whether a strategy profile of an iterated EBG is a Nash equilibrium, we will see how we can easily tackle derived decision problems for engineering Electric Boolean Games.

A resource redistribution for an EBG $\mathcal{B} = (N, \Sigma, \Phi, c, e)$ is an endowment function $e' : N \rightarrow \mathbb{N}$ such that

$$\sum_{i \in N} e(i) = \sum_{i \in N} e'(i).$$

Remark 14. *Let an EBG $\mathcal{B} = (N, \Sigma, \Phi, c, e)$. There is finite number of resource redistributions for \mathcal{B} .*

In [19], the authors studied the problems of determining whether there is a resource redistribution such that a strategy profile is a Nash Equilibrium (rational construction), and of determining whether there is a resource redistribution such that a strategy profile is not a Nash Equilibrium (rational elimination). For the iterated setting we propose the following decision problems.

Definition 15 (Construction and elimination). *Let \mathcal{B} be an electric boolean game, and σ be a finite memory strategy profile. The Rational Construction (RC) problem asks whether there is a resource redistribution such that σ is a Nash equilibrium. The Rational Elimination (RE) problem asks whether there is a resource redistribution such that σ is not a Nash equilibrium.*

Theorem 16. *The RC problem and the RE problem are PSPACE-complete.*

The non-deterministic procedures outlined in the proof of Theorem 16 are sufficient to characterise an optimal upper-bound of the problems. In the case of RE, there exists a more practical deterministic algorithm. Indeed, the result of [19, Corr. 4] carries over in the iterated setting.

Proposition 17. *Let an endowment e be given. The endowment e^i is the resource redistribution of e such that all resources are allocated to player i . The strategy profile σ is eliminable in $\mathcal{B}^{c,e}$ iff for some player i , $\sigma \notin \text{NE}(\mathcal{B}^{c,e^i})$.*

This hints at a “more practical” algorithm to solve RE: for each player i , test whether $\sigma \notin \text{NE}(\mathcal{B}^{c,e^i})$. Return “yes” as soon as a test succeeds. Return “no” when all $|N|$ tests failed.

5 Conclusion

In this paper we presented a preliminary result on the Electric Boolean Games introduced in [19]. We considered the iterated setting where the objectives are specified as LTL formulas. We showed the PSPACE-completeness of Nash equilibrium membership, thus matching the complexity bounds of [17] for the non quantitative setting of iterated Boolean Games. In order to establish this result, we extended existing techniques for plain LTL to an extension of LTL with electric constraints. This result is used to characterise the complexity of two problems of resource redistribution that can serve at social-welfare engineering.

As future research direction, we plan to investigate the Nash equilibrium non-emptiness and Nash equilibrium synthesis. We believe that Construction 2 can be extended in order to construct a concurrent game with the property that it contains a pure Nash equilibrium iff the electric boolean game does. To the best of our knowledge, the obtained class of concurrent games is rather novel and has yet to be studied.

References

- [1] Shaull Almagor, Guy Avni & Orna Kupferman (2015): *Repairing Multi-Player Games*. In: *CONCUR 2015, LIPIcs* 42, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 325–339, doi:10.4230/LIPIcs.CONCUR.2015.325.
- [2] Rajeev Alur, Thomas A. Henzinger, Freddy Y. C. Mang, Shaz Qadeer, Sriram K. Rajamani & Serdar Tasiran (1998): *MOCHA: Modularity in Model Checking*. In: *CAV 1998*, Springer, pp. 521–525, doi:10.1007/BFb0028774.
- [3] Eugene Asarin, Oded Maler & Amir Pnueli (1994): *Symbolic Controller Synthesis for Discrete and Timed Systems*. In: *Hybrid Systems II*, pp. 1–20, doi:10.1007/3-540-60472-3_1.
- [4] Christel Baier & Joost-Pieter Katoen (2008): *Principles of Model Checking*. The MIT Press.
- [5] Elise Bonzon, Caroline Devred & Marie-Christine Lagasquie-Schiex (2010): *Argumentation and CP-Boolean Games*. *International Journal on Artificial Intelligence Tools* 19(4), pp. 487–510, doi:10.1142/S0218213010000297.
- [6] Elise Bonzon, Marie-Christine Lagasquie-Schiex & Jérôme Lang (2009): *Dependencies between players in Boolean games*. *International Journal of Approximate Reasoning* 50(6), pp. 899–914, doi:10.1016/j.ijar.2009.02.008.
- [7] Elise Bonzon, Marie-Christine Lagasquie-Schiex, Jérôme Lang & Bruno Zanuttini (2006): *Boolean Games Revisited*. In: *ECAI 2006, Frontiers in Artificial Intelligence and Applications* 141, IOS Press, pp. 265–269.
- [8] Elise Bonzon, Marie-Christine Lagasquie-Schiex, Jérôme Lang & Bruno Zanuttini (2009): *Compact preference representation and Boolean games*. *Autonomous Agents and Multi-Agent Systems* 18(1), pp. 1–35, doi:10.1007/s10458-008-9040-2.
- [9] Tomás Brázdil, Krishnendu Chatterjee, Vojtech Forejt & Antonín Kucera (2015): *MultiGain: A Controller Synthesis Tool for MDPs with Multiple Mean-Payoff Objectives*. In: *TACAS 2015*, pp. 181–187, doi:10.1007/978-3-662-46681-0_12.
- [10] Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur & Mathieu Sassolas (2016): *Non-Zero Sum Games for Reactive Synthesis*. In: *LATA 2016*, pp. 3–23, doi:10.1007/978-3-319-30000-9_1.
- [11] Thomas Brihaye, Julie De Pril & Sven Schewe (2013): *Multiplayer Cost Games with Simple Nash Equilibria*. In: *LFCS 2013*, pp. 59–73, doi:10.1007/978-3-642-35722-0_5.
- [12] Krishnendu Chatterjee & Laurent Doyen (2012): *Energy parity games*. *Theor. Comput. Sci.* 458, pp. 49–60, doi:10.1016/j.tcs.2012.07.038.
- [13] Krishnendu Chatterjee, Thomas A. Henzinger & Marcin Jurdzinski (2005): *Mean-Payoff Parity Games*. In: *(LICS 2005)*, pp. 178–187, doi:10.1109/LICS.2005.26.
- [14] Krishnendu Chatterjee, Mickael Randour & Jean-François Raskin (2014): *Strategy synthesis for multi-dimensional quantitative objectives*. *Acta Informatica* 51(3-4), pp. 129–163, doi:10.1007/s00236-013-0182-6.
- [15] Paul E. Dunne, Wiebe van der Hoek, Sarit Kraus & Michael Wooldridge (2008): *Cooperative Boolean games*. In: *AAMAS 2008, IFAAMAS*, pp. 1015–1022.
- [16] John Grant, Sarit Kraus, Michael Wooldridge & Inon Zuckerman (2011): *Manipulating Boolean Games through Communication*. In: *IJCAI 2011, IJCAI/AAAI*, pp. 210–215, doi:10.5591/978-1-57735-516-8/IJCAI11-046.
- [17] Julian Gutierrez, Paul Harrenstein & Michael Wooldridge (2015): *Iterated Boolean games*. *Information and Computation* 242, pp. 53–79, doi:10.1016/j.ic.2015.03.011.
- [18] Paul Harrenstein (2004): *Logic in conflict*. Ph.D. thesis, Utrecht University.
- [19] Paul Harrenstein, Paolo Turrini & Michael Wooldridge (2015): *Electric Boolean Games: Redistribution Schemes for Resource-Bounded Agents*. In: *AAMAS 2015, ACM*, pp. 655–663.

- [20] S. Rao Kosaraju & Gregory F. Sullivan (1988): *Detecting Cycles in Dynamic Graphs in Polynomial Time (Preliminary Version)*. In: *STOC 1988*, ACM, pp. 398–406, doi:10.1145/62212.62251.
- [21] Alessio Lomuscio, Hongyang Qu & Franco Raimondi (2009): *MCMAS: A Model Checker for the Verification of Multi-Agent Systems*. In: *CAV 2009*, Springer, pp. 682–688, doi:10.1007/s10009-015-0378-x.
- [22] Youssouf Oualhadj & Nicolas Troquard (2016): *Rational verification in Iterated Electric Boolean Games*. CoRR abs/1604.03773. Available at <http://arxiv.org/abs/1604.03773v1>.
- [23] Stavros Tripakis & Karine Altisen (1999): *On-the-Fly Controller Synthesis for Discrete and Dense-Time Systems*. In: *FM'99*, pp. 233–252, doi:10.1007/3-540-48119-2_15.
- [24] Michael Wooldridge, Ulle Endriss, Sarit Kraus & Jérôme Lang (2013): *Incentive engineering for Boolean games*. *Artificial Intelligence* 195, pp. 418 – 439, doi:10.1016/j.artint.2012.11.003.